

Evaluating the End-User Experience of Private Browsing Mode

Ruba Abu-Salma

University College London (UCL)

Benjamin Livshits

Brave Software / Imperial College London

ABSTRACT

In this paper, we investigate why users of private browsing mode misunderstand the benefits and limitations of private browsing. We design and conduct a three-part study: (1) an analytic evaluation of the user interface of private mode in different browsers; (2) a qualitative user study to explore user mental models of private browsing; (3) a participatory design study to investigate why existing browser disclosures, the in-browser explanations of private mode, do not communicate the actual protection of private mode.

We find the user interface of private mode in different browsers violated well-established design guidelines and heuristics. Further, most participants had incorrect mental models of private browsing, influencing their understanding and usage of private mode. We also find existing browser disclosures did not explain the primary security goal of private mode. Drawing from the results of our study, we extract a set of recommendations to improve the design of disclosures.

Author Keywords

Human-Computer Interaction; Usable Security and Privacy.

CCS Concepts

•Security and privacy → Social aspects of security and privacy; Usability in security and privacy;

INTRODUCTION

Prior work has extensively explored users' online privacy concerns when using the Internet [1,5,20,31,35–38]. For example, a survey of 1,002 US respondents (conducted by the Pew Research Center in 2013) found that respondents were concerned about their personal information being available online [38]. Respondents also felt strongly about controlling who had access to their behavioural data and communications, including family members, partners, friends, employers, advertisers, and government agencies. In 2015, Angulo and Ortlieb conducted a user study to investigate users' concerns with regards to "online privacy-related panic" incidents [5]. They found that online tracking, reputation loss, and financial harm were the most frequently reported incidents by participants.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CHI '20, April 25–30, 2020, Honolulu, HI, USA.

© 2020 Association for Computing Machinery.

ACM ISBN 978-1-4503-6708-0/20/04 ...\$15.00.

<http://dx.doi.org/10.1145/3313831.3376440>

Prior work has also found that users are willing to take measures to protect their online privacy. In the same Pew Research Center survey [38], a clear majority (86%) of respondents reported they had taken steps to remove or hide their "digital footprints," including clearing their browsing history and cookies. Further, Kang et al. conducted a user study to investigate how users would react to security and privacy risks [28]; 77% of non-technical participants reported taking several measures to protect their "digital traces," including the use of private browsing mode.

As we can see, users have serious concerns about their online privacy, and try to employ different strategies or use different privacy-enhancing tools to protect it. In this work, we focus on evaluating the end-user experience of one of these tools: **private browsing mode**¹. Private browsing is a privacy-enhancing technology that allows a user to browse the Internet without saving information (e.g., browsing history, cookies, temporary files) about the websites they visited in private mode on their local device [2]. As of today, all major web browsers have a private browsing mode.

The primary security goal of private browsing is that a local attacker – such as a family member, a friend, or a work colleague – who takes (physical or remote) control of the user's machine *after* the user exits a private browsing session should find no evidence of the websites the user visited in that session [2]. That is, a local attacker who has access to the user's machine at time T should learn nothing about the user's private browsing activities prior to time T . Therefore, private browsing does not protect against a local attacker who controls the user's machine *before* or *during* a private browsing session.

Further, private browsing does not aim to protect against a web attacker who, unlike a local attacker, does not control the user's machine but controls the websites visited by the user in private mode [2]. It also does not hide private browsing activities from the browser vendor, Internet service provider (ISP), employer, or government.

Previous user studies have quantitatively – mainly through survey studies – investigated whether users are aware of private browsing, what they use it for, and whether they understand what protection it provides [7, 13, 22, 26, 32, 48]. However, these studies have not investigated *why* most users misunderstand the benefits and limitations of private browsing mode. Further, many of the recruited participants in these studies were unaware of or had not used private mode. In this work, we address these research gaps by designing and conducting

¹ In this paper, we use the terms "private browsing mode," "private browsing," and "private mode" interchangeably.

a three-part study, where we recruited 25 demographically-diverse participants (**both users and non-users of private mode**) for the second and third parts of the study.

First, we use an analytic approach combining cognitive walk-through and heuristic evaluation to inspect the user interface of private mode in different web browsers. Second, we conduct a qualitative, interview-based study to explore user mental models of private browsing and its security goals. Third, we perform a participatory design study to investigate whether existing browser disclosures, the full-page explanations browsers present when users open a new private tab or window in private mode, communicate the security goals of private browsing to users. We ask participants to critique the disclosures of Brave, Chrome, and Firefox, and then design new ones.

We summarize our key findings below:

- We identify usability issues in the user interface of private mode in different browsers. We find some of these issues hampered the adoption of private mode (e.g., minimal feedback, use of technical jargon, lengthy disclosures).
- We find participants had inappropriate mental models of private mode. Further, almost all participants did not understand the security goal of private browsing. For example, some participants sent emails unencrypted in private mode, incorrectly believing private mode achieved confidentiality.
- We find most participants who used private mode performed their private browsing activities while being authenticated to their personal online account (mainly their Google account), incorrectly believing their browsing or search history would get deleted from Google's records after exiting private mode.
- We find none of the three studied browser disclosures communicated the security goal of private mode. Our participants also pointed out that disclosures did not explain where information related to a private browsing session would get deleted from, and when.

Drawing from these findings, we extract a set of guidelines to improve the design of disclosures. We also propose a new disclosure design (see Section 5: 'Discussion').

RELATED WORK

User Studies of Private Browsing Mode

Prior work has **quantitatively** (mainly through survey studies) investigated whether users are aware of private browsing, what they use it for, and whether they understand what protection it provides. In [22], Gao et al. conducted a survey of 200 Mechanical Turk (MTurk) respondents in the US, examining their private browsing habits. They found that one-third of respondents were not aware of private browsing. Those who had used private browsing reported using it for protecting personal information, online shopping, or visiting "embarrassing websites." Gao et al. concluded that "browser designers [should think of] various ways to [better] inform users."

In 2017, DuckDuckGo, an Internet search engine, surveyed a sample of 5,710 US respondents, recruited via SurveyMonkey [13]. Respondents were asked to share their experience

with private browsing. Again, one-third of respondents reported they had not heard of private browsing. Of those who had used private browsing, one-third used it frequently, and three-quarters were not able to accurately identify the benefits of private browsing.

Using a similar study to [13], Bursztein ran an online survey of 200 US respondents (via Google Consumer Surveys) in 2017 [7]. He found about one-third of surveyed respondents did not know about private browsing. Of those who were aware of the technology, only 20% had used it. Further, about one-half preferred not to disclose what they used private browsing for. Bursztein concluded that "surveys are clearly not the best approach to understand why people are using the private browsing mode because of the embarrassment factor."

Recently, Wu et al. surveyed 460 US respondents through MTurk [48]. Respondents were randomly assigned one of 13 different browser disclosures related to private mode. Based on the disclosure they saw, respondents were asked to answer a set of questions to assess their understanding of private mode. Wu et al. found that existing disclosures did not inform users of the benefits and limitations of private mode. They concluded that disclosures "should be redesigned."

Habib et al. conducted a user study to observe the private browsing habits of over 450 US participants using software monitoring [26]. They then asked participants to answer a follow-up survey (using MTurk) to investigate discrepancies, if any, between observed and self-reported private browsing habits. They found that the primary use cases of private mode were consistent across observed and self-reported data. They also found that most participants overestimated the benefits of private mode, concluding by supporting "changes to private browsing disclosures."

Summary. Prior work has employed **quantitative** methods to investigate users' private browsing habits. However, prior work has not investigated **why** users misunderstand the benefits and limitations of private browsing. Further, many of the recruited participants in prior user studies were unaware of or had not used private mode. In this work, we address these research gaps by designing and conducting a three-part user study: (1) a usability inspection of private mode in different web browsers, (2) a qualitative, interview-based user study, and (3) a participatory design study. We also recruit both users and non-users of private mode.

Security and Privacy Design

Within web browsers, prior work has investigated the design of alert messages and warnings [3, 4, 12, 14–16, 18, 42, 43], browser security indicators [17, 21, 39], site trustworthiness [9, 34], privacy policies [45, 47], storage policies [46], and ad personalization [30]. However, prior work has heavily focused on the design of warning messages – especially phishing warnings [3, 12, 14, 15] and SSL warnings [3, 16, 18, 42, 43] – in order to capture users' attention, improve their comprehension, and warn them away from danger. For example, Egelman et al. recommended that phishing warning messages should be active (i.e., interrupt the user flow) and should be distinguishable by severity [14]. They also suggested it should be difficult

for users to click-through phishing warnings, by requiring users to bypass several screens in an attempt to dissuade users from ignoring warnings. Additionally, Egelman and Schechter showed that changes to the look and feel of phishing warnings have resulted in more users noticing them [15]. Felt et al. recommended warning designers use opinionated design to improve user adherence to warnings [16].

Summary. The aforementioned work has focused on the design of browser security warnings to improve their efficacy. However, our study focuses on designing browser disclosures that inform users of the benefits and limitations of a privacy tool (private mode). Although we draw inspiration from prior work, **we answer a different important question of how to design disclosures to help users appropriately use private mode.** We do so by employing **participatory design** [40]. Unlike warning designers who have explored different ideas – such as changing the design of a warning or using attractors [6] – to improve user attention to and comprehension of warnings, we choose, in this work, to engage users in the design of disclosures related to private mode.

METHODOLOGY

We conducted a three-part study: (1) a hybrid analytic approach combining cognitive walkthrough and heuristic evaluation; (2) a qualitative, interview-based user study; (3) a participatory design study. Our study was reviewed and approved by our organization’s ethics committee.

For the second and third parts of our study, we first conducted five unstructured (open-ended) face-to-face interviews, lasting for 60 minutes on average each. The emerging themes from these five interviews helped us design the study script we used to conduct our main interviews: 25 semi-structured face-to-face interviews lasting for 90 minutes on average each (see Table 1 in Section 4: ‘Results’). We describe the script on the next page (see Part 2 and Part 3).

Research Questions

In this paper, we answer the following research questions:

- **RQ1:** Does private mode in different web browsers suffer from poor usability that hampers the widespread adoption and use of private browsing?
- **RQ2:** How do users perceive the term “private browsing?”
- **RQ3:** What are users’ mental models of private browsing (as a privacy-enhancing technology) and its security goals?
- **RQ4:** How do users perceive those who use private browsing? Do users perceive the routine use of private browsing as “paranoid” or “unnecessary?”
- **RQ5:** How do users’ mental models and perceptions influence their usage of private browsing?
- **RQ6:** Why do existing browser disclosures (related to private browsing) misinform users of the benefits and limitations of private browsing?
- **RQ7:** How can the design of disclosures be improved?

Recruitment

To recruit our participants (for the second and third parts of the study), we posted flyers and distributed leaflets in London (UK). We asked interested participants to complete an online screening questionnaire², which about 500 completed. We aimed to recruit a demographically-diverse sample of participants. Hence, we included a number of demographic questions about gender, age, race, educational level, and employment status. We also assessed participants’ technical knowledge; we considered participants as technical if two out of three of the following were true [44]: (1) participants had an education in, and/or worked in, the field of computer science, computer engineering, or IT; (2) they were familiar with or an expert in at least one programming language (e.g., C++); (3) people usually asked them for computer-related advice. Further, we provided participants with a list of different web browsers, and then asked which browsers they used, what they used each browser for (in case they used multiple browsers), which browser they used the most, and how many hours they spent daily on their desktop and mobile phone browsing.

Additionally, we asked participants to list the digital security requirements they had at school or work, how often they received cybersecurity training, and whether they felt at risk due to their school work or job duties. In [23], Gaw et al. found that people perceived the “universal, routine use of encryption as paranoid.” In this work, we aimed to explore whether participants perceived the use of private mode as paranoid.

Part 1: Usability Inspection

Usability inspection is aimed at finding usability problems in the user interface design. Unlike empirical user studies (see Part 2 and Part 3 on the next page), a user interface is inspected by evaluators without engaging users (i.e., without recruiting participants to assess the usability of a system). Although it is important to bring users into the design process, evaluating a design without users can also provide benefits [29].

To answer RQ1, we used a hybrid approach combining cognitive walkthrough and heuristic evaluation to inspect the user interface of private mode in five different web browsers: Brave, Chrome, Internet Explorer, Firefox, and Safari³. We:

1. provided a detailed description of the user interface.
2. defined the users and their goals.
3. defined the tasks the users would attempt (e.g., accessing a web page in private mode).
4. divided each task into a sequence of sub-tasks or actions (e.g., selecting the “New Private Window” option).
5. walked through each task workflow step-by-step through the lens of the users (e.g., what they would look for).
6. looked for and identified usability problems (for each action) based on the ten heuristics described in [33].
7. specified where the usability problem was in the user interface, how severe it was, and possible design fixes.

² <https://tinyurl.com/chi2020-1>

³ <https://tinyurl.com/chi2020-2>

Part 2: Interview-Based Study

After inspecting the user interface of private mode, we aimed to answer RQ2–RQ5 by qualitatively investigating participants' mental models of private browsing and its security goals, as well as exploring how participants perceived those who (regularly or occasionally) used private browsing. We also aimed to understand how participants' mental models and perceptions influenced their understanding and usage of private mode. Hence, we explored the following themes:

Mental models of "private browsing." We asked participants whether they had heard of the term "private browsing," and, if so, whether they felt confident explaining what it meant. We then asked them to explain what it meant to browse privately. We provided participants with a large pad of paper and a 24-color pack of markers, giving them the option to draw their mental models of private browsing. We also asked participants to describe the benefits and drawbacks of browsing privately.

By asking these questions, we aimed to investigate participants' conceptual understanding of the term "private browsing," and how this understanding influenced their mental models and usage of private mode (as a privacy-enhancing technology), as we describe next in detail.

Mental models of private mode (as a privacy tool). After exploring participants' general mental models of the term "private browsing," we asked participants whether they had browsed in private mode and, if so, whether they felt confident explaining what it meant to open a private tab or window. We then asked them to explain the difference, if any, between default (non-private) mode and private mode.

We also aimed to understand how participants perceived the security goals of private mode. Hence, we asked participants about the entities that could learn about their private browsing activities (e.g., visited websites in private mode), and how. We wanted to explore whether participants understood the primary goal of private mode: protecting against a local attacker who controls a user's machine after the user exits private mode.

Perceptions of users of private mode. We then asked participants to explain how they perceived the use of private mode. We aimed to investigate whether participants perceived users of private mode as paranoid.

Expectations. We asked participants to describe what they would expect from private mode. We also investigated whether participants' familiarity with private mode affected the robustness of their mental models. Therefore, we asked participants to list the web browsers that they used (as well as those they did not necessarily use) and that they considered having a private mode that met their expectations.

Private browsing usage. Finally, we aimed to explore how participants' mental models influenced their usage of private mode. Hence, we asked participants who used, or had used in the past, private mode to share their private browsing habits. We asked them what they used private mode for, how often they used it, and where they used it. We also asked them to explain what they liked and disliked about private mode.

Part 3: Participatory Design Study

After exploring our participants' mental models and usage of private mode, we aimed to investigate why browser disclosures (related to private browsing) did not communicate the actual benefits and limitations of private browsing. We also sought to improve the design of existing browser disclosures. Hence, we performed a participatory design study to solicit new disclosure designs from our participants.

Assessing participants' knowledge of private mode (before tutorial). To answer RQ6 and RQ7, we asked participants to take a short quiz⁴ to further test their knowledge of private mode. We asked them to answer seven questions about a private mode that worked properly.

We also asked participants whether they were familiar with the following items that appeared on almost all existing browser disclosures, and whether they felt confident explaining what each item meant: browsing history file, cookies, search items, bookmarks, downloads, and temporary files.

Giving a tutorial. We then gave participants a 15-minute tutorial, explaining the primary security goal of private browsing, the difference between default (non-private) mode and private mode, and why private browsing did not protect against website fingerprinting and, hence, website tracking and ad targeting. Further, we explained the different items/files that most browsers claimed to delete when a user exited private mode. We also explained the different privacy features that had been recently added by some web browsers (e.g., Brave's Private Tabs with Tor). Finally, we explained the difference between a private tab, a private window, and a private session.

Assessing participants' knowledge of private mode (after tutorial). To evaluate whether participants' knowledge of private browsing had improved after the tutorial, we asked participants to take the same quiz we gave them previously. However, we shuffled the questions to minimize bias.

Critiquing existing disclosures. We then asked each participant to critique the browser disclosures of three browsers: Brave, Chrome, and Firefox. To minimize bias, disclosures were assigned to each participant randomly. We chose these three disclosures because Chrome and Firefox were the most frequently-used browsers by participants, whereas Brave was launched with privacy as a key selling point.

We conducted a within-subject study; we showed participants one disclosure at a time. We then asked them to describe what they felt about the disclosure, how useful they felt the explanation was, what about the explanation would make them decide to use private mode, and what else they would like the disclosure to tell them. We then showed participants the second disclosure and followed-up by asking the same questions we asked about the first disclosure they saw. We also asked participants to compare the second disclosure to the first one, and then explain whether they would be more or less likely to use private mode if they saw this disclosure or the prior one. We then showed participants the third disclosure and asked them the same questions we previously asked.

⁴ <https://tinyurl.com/chi2020-3>




Figure 1: Brave disclosure




Figure 2: Chrome disclosure




Figure 3: Firefox disclosure

Soliciting new disclosure designs. We then performed a participatory design study to solicit new disclosure designs from our participants. We asked participants to describe private browsing as if they were explaining it to someone new to this privacy-enhancing technology. We prompted our participants as follows: “We would like you to design a browser disclosure that clearly explains the benefits and limitations of private browsing. While designing, think about what would make you use private mode, what information you would want to know, what information you would want to omit, and how you would want the disclosure to look.” We gave participants a large pad of paper and a 24-color pack of markers to design their disclosures, giving them the option to draw.

Finally, we asked participants to share their thoughts on the following names: “Private Browsing,” “InPrivate Browsing,” and “Incognito Browsing,” and suggest a new name, if any.

Pilot Study

Quiz piloting. After developing our quiz (see Part 3 in ‘Methodology’), we conducted cognitive interviews with five participants to test our quiz. Cognitive interviewing is a method used to test questionnaires to glean insights into how participants might interpret questions [25].

Main study piloting. To pre-test the second and third parts of our study (pre-screening questionnaire, study script, and quiz), we conducted a small-scale pilot study of five semi-structured interviews (using convenience sampling [25]). Additionally, we asked ten computer security and privacy researchers and experts to review the study. We used the findings to identify potential problems (e.g., time, cost, adverse events) in advance prior to conducting the full-scale study.

Data Analysis

Usability inspection (Part 1). Two expert HCI researchers inspected the user interface of private mode in Brave, Chrome, Internet Explorer, Firefox, and Safari. We did so independently before discussing the findings and aggregating all the uncovered issues in a larger set.

Interview-based and participatory design studies (Part 2 and Part 3). We conducted, transcribed, and analyzed all five unstructured and 25 semi-structured interviews (the study’s main interviews). We observed data saturation [11, 24, 41] between the 20th and the 25th semi-structured interview; i.e., no new codes emerged in interviews 20–25, and, hence, we stopped recruiting participants. Data saturation is commonly taken to

indicate, on the basis of the data that has been collected and analyzed, further data collection and analysis are unnecessary.

Two researchers independently coded all interview transcripts and image data using Grounded Theory [11, 24, 41]. The researchers created two codebooks: one for the interview transcripts and one for the image data. After creating the final codebooks, we tested for the inter-rater reliability. The average Cohen’s kappa coefficient (κ) for all themes in the interview transcripts and image data was 0.77 and 0.89, respectively. A κ value above 0.75 is considered excellent agreement [10, 19].

RESULTS

In this section, we present the results of our study.

Demographics

Table 1 summarizes the demographics of our sample ($n=25$). We interviewed 11 male, 13 female, and one non-binary participants. Participants’ ages ranged from 18 to 75. 13 identified as white, four as black, four as Asian, two as Hispanic, and two as mixed-race. Eight reported having a college (or an undergraduate) degree, and nine a graduate (or postgraduate) degree. Two reported having secondary education, and four some post-secondary education (i.e., some college education without a degree). Two participants mentioned having vocational training (VOC). Nine participants were either high-school or university students, 12 employed, two unemployed, and one retired. One participant preferred not to indicate their employment status. According to the definition we used to assess our participants’ technical knowledge (see Section 3.2: ‘Methodology’), 17 qualified as technical.

Our participants used a wide range of web browsers (both on desktop/laptop and mobile phone). Chrome was the most used browser by participants, followed by Safari, Firefox, Internet Explorer, and Brave, respectively. Three participants (P01; P03; P25) used the Tor browser.

Participants daily spent between five and 17 hours browsing the Internet. Desktop/laptop browsing overtook smartphone surfing, with the exception of three participants (P02; P12; P16). Further, most participants (22 out of 25) used multiple browsers. For example, 21 reported they used one browser for social activities and one for work-related activities.

19 participants reported they used (or had used in the past) private mode. Three (P12; P16; P24) were aware of private mode, but had not browsed in it. Three (P02; P11; P23) did not know private mode existed.

	Gender	Age	Race	Education	Employment
P01	Male	25–34	White	Ph.D.	Student
P02	Male	45–54	Mixed race	B.A.	Unemployed
P03	Male	45–54	White	Ph.D.	Unemployed
P04	Female	18–24	Black	High-school	Student
P05	Female	25–34	White	B.A.	Employed
P06	Male	35–44	White	M.Sc.	Employed
P07	Female	18–24	White	B.A.	Employed
P08	Female	25–34	Asian	High-school	Student
P09	Male	18–24	Asian	M.Sc.	Employed
P10	Male	25–34	White	Some college	Employed
P11	Female	25–34	White	M.Sc.	Employed
P12	Female	45–54	White	Some college	Employed
P13	Male	25–34	Mixed race	B.A.	Employed
P14	Male	18–24	Hispanic	B.A.	Employed
P15	Female	25–34	Asian	B.Sc.	Other
P16	Female	45–54	Black	VOC	Employed
P17	Female	18–24	White	Ph.D.	Student
P18	Non-binary	35–44	White	M.Sc.	Employed
P19	Female	35–44	Black	B.Sc.	Self-employed
P20	Male	18–24	White	Some college	Retired
P21	Male	25–34	White	VOC	Student
P22	Male	18–24	Asian	Ph.D.	Student
P23	Female	25–34	White	M.Sc.	Student
P24	Female	25–34	Black	B.Sc.	Student
P25	Female	65–74	Hispanic	Some college	Student

Table 1: Semi-structured interview participant demographics.

Part 1: Usability Inspection

We used an analytic approach combining cognitive walk-through and heuristic evaluation to inspect the user interface of private mode in five different web browsers (desktop versions). Our findings are as follows:

Public mode as the default mode. In all browsers (including the ones we inspected), the default mode is the public one. To browse in private mode, users need to select (from a hidden drop-down list) “New Incognito window” in Brave and Chrome, or “New private window” in Internet Explorer, Firefox, and Safari. We hypothesize (and find in Section 4.3: ‘Part 2: Interview-Based Study’) most users are unaware of the hidden list, which explains why most users do not know about private mode. This violates Nielsen’s heuristics of *visibility of system status* [27] and *aesthetic and minimalist design* [27].

Multiple windows and tabs. Users cannot open a private tab in a public window, and vice-versa; that is, users can only open public (private) tabs in public (private) windows – which we regard as good design. Further, users can only re-open the most recently-closed public tabs, and not private ones.

Although users can open multiple public and private windows, feedback is minimal. For example, in Safari, when users enter private mode, there is no appropriate feedback – through the user interface – that communicates to users that they are currently browsing in private mode. There is only a short line of text (using a small font size) at the top of the page that says: “Private Browsing Enabled.” In Brave and Firefox, the background changes from white to purple. Both browsers do not explain why purple was chosen by browser designers to distinguish between public and private modes.

Use of jargon. Both Brave and Chrome refer to private mode as “Incognito window,” and Internet Explorer, Firefox, and Safari as “private window.” This violates Nielsen’s heuristic of *match between the system and the real world* [27], making the assumption that users’ understanding and interpretation of words or terms would be the same as browser designers and developers. We also hypothesize that users would build their own mental models of private mode when encountering these

terms, which could strongly impact how users would perceive and use the mode in real life. We explore user mental models of private mode in the following section (Part 2).

Wordy disclosures. When users enter private mode, a browser disclosure is shown to them. The disclosure is meant to explain the benefits and limitations of private browsing. However, the disclosures of all inspected browsers (except that of Firefox) are lengthy and full of jargon, violating Nielsen’s heuristic of *match between the system and the real world* [27]. Further, browser disclosures do not explain the primary security goal of private mode. In Firefox, the disclosure is relatively short, but, also, does not explain the security goal of private mode.

Further, in all five browsers, users are presented with these disclosures only once (when they open a private window or tab), violating Nielsen’s heuristics of *recognition rather than recall* [27] and *help and documentation* [27].

In Section 4.4 (‘Part 3: Participatory Design Study’), we present the results from our participants who critiqued three existing browser disclosures and suggested several design options for improvement.

Private browsing and Tor. Brave has recently added Tor to its private windows. Brave users can now open a “New window,” “New Incognito window,” or “New private window with Tor.” Both Incognito windows and private windows with Tor have the same purple background and lengthy disclosures, which could lead users to browse in one instead of the other, violating Nielsen’s heuristic of *visibility of system status* [27]. Further, the browser disclosures of both windows do not clearly explain how private mode and Tor are two different privacy tools.

Part 2: Interview-Based Study

The main purpose of qualitative research is to explore a phenomenon in depth, and not to investigate whether or not findings are statistically significant or due to chance [25]. Although we report how many participants mentioned each finding as an indication of prevalence, our findings are not quantitative. Further, a participant failing to mention a particular finding does not imply they disagreed with that finding; they might have failed to mention it due to, for example, recall bias [25]. Thus, as with all qualitative data, our findings are not necessarily generalizable beyond our sample. However, they suggest several future research avenues, and can be later supplemented by quantitative data.

In this section and the next section, we present the results of the second and third parts of the study (n=25).

Mental models of “private browsing.” We aimed to investigate our participants’ conceptual understanding of the term “private browsing.” 18 out of 25 (a clear majority) had heard of the term, and 17 felt confident explaining what the term meant⁵. 16 out of 17 were users of (or had used in the past) private mode. One participant (P11) was a non-user.

⁵ It is worth to mention that only three out of those 17 participants associated the term “private browsing” with private mode. We speculate that this was because the three participants used – regularly or occasionally – private mode.

We then asked all participants to explain what “private browsing” meant to them. Five out of 25 associated the term with private browsing mode, mentioning the following: “the window that has a man with a coat and a pair of eye glasses” (x4); “going undercover or incognito” (P04). All five participants were referring to the “Incognito window” in Google Chrome. Further, five participants thought of the term in connection with network-encrypted communications or secure browser connections (i.e., webpages running HTTPs), three with end-to-end encrypted communications, three with anonymous communications (using Tor or VPN), and three with user authentication (both one-factor and two-factor authentication). One participant (P17) associated “private browsing” with both network encryption and authentication. Additionally, P15 described the term as the ability to browse the Internet “without getting infected with a virus.”

Further, eight participants mentioned the terms “privacy” and “online privacy” to explain what “private browsing” meant to them: P01–P05, P07, and P12–P14 defined the term as having control over how users’ online information was handled and shared. P09, P20, P22, and P24 referred to the term as the ability to manage and “regulate” one’s social space.

Mental models and usage of private mode (as a privacy tool). After exploring our participants’ conceptual understanding of the term “private browsing,” we aimed to investigate how this understanding influenced participants’ mental models and usage of private mode (as a privacy tool). We identified three types of users of private mode: regular users, occasional users, and former users. We explain each type as follows:

1. Regular users: Two participants (P01 and P17) were regular users of private mode. They performed all their browsing activities in private mode. They described themselves as “paranoid” and “cautious.” P01 mentioned that the routine use of private mode made them feel “safer” and “more comfortable.” Further, P01 used Safari’s private mode to protect against shoulder-surfing. They explained that Safari did not have a visual user interface element that indicated a user was currently browsing privately. However, when probed, P01 (as well as P17) did not know that staying in private mode for a long period of time could easily enable fingerprinting and, hence, website tracking (a threat that both participants thought they were protected against by regularly browsing in private mode).

2. Occasional users: Out of 25, 15 participants used private mode occasionally depending on their browsing activities and the websites they visited. They did not necessarily use the mode to visit “embarrassing websites.” Many used private mode for online shopping (e.g., purchasing a surprise gift for a family member or a friend), logging into an online service using a different account, and/or debugging software.

3. Former users: Two participants (P13 and P19) reported they had used private mode before, but they stopped using it for the following reasons:

- *Lack of usability.* P13 and P19 mentioned that entries added to the history file would get deleted if they exited private mode, negatively impacting user experience. P13 also mentioned that private mode was “useless” because users could

delete information about websites visited in default mode by manually clearing their browsing history file (a view shared by P12 and P16).

- *Lack of utility.* P13 stopped using private mode because they thought that web browsers did not allow extensions to run in private mode (although users could manually enable extensions in private mode in most browsers). This finding was also shared by five other participants.
- *Misconceptions about private mode.* P13 perceived those who used private mode as people who “had something to hide” or “were up to no good,” influencing P13’s decision to stop using private mode; P13 did not want to be perceived by others as a “cybercriminal.” Some participants shared this perception, as we discuss later in this section.

17 out of 25 participants reported they mainly used private mode in public spaces using shared devices, mainly coffee shops, libraries, and airports. They performed browsing activities they regarded as sensitive in private mode. For example,

“I usually use Incognito in ... you know ... in Google when I work at [coffee shop] because I connect to the Internet using insecure or public Wi-Fi. My laptop consistently warns me. So, I use Incognito to encrypt my data and hide it from people around me ... Better to be safe!” (P05)

“I usually use the public or ... shared workstations in my school’s library. You don’t need to login because there is one account shared by all students. I open a private tab or window to download files that I want to be removed after I close the browser ... By the way, I also use a private window to send an encrypted email.” (P17)

Surprisingly, P17 was a regular user of Safari that locally deleted files downloaded in its private mode. However, P17 did not notice he was using Firefox on the library’s computer, which did not delete private browsing downloads.

“I make a bank transfer or access my personal accounts – you know, like Facebook – when I use one of the computers that all passengers can use ... I am talking about the computers you find in an airport lounge ... I open a private window.” (P07)

“I use Incognito to search for new jobs. Hmm, I do not want my boss or company to know ...” (P18)

“If I do not have Tor installed, I will use Incognito.” (P09)

We also found six participants who tended to use private mode to visit malicious webpages. For example,

“I sometimes encounter a message that warns me from accessing a bad webpage. I usually ignore the warning and open the page in a private window ... Feels safer!” (P14)

Alarmingly, we found *all* participants who used or had used private mode (x19) browsed privately while being authenticated to their Google or YouTube account, incorrectly believing their search history would get deleted from Google or YouTube’s records after exiting private mode.

Additionally, we found that some participants (11 out of 25) perceived those who used private mode as people who “cared

about their online privacy,” “had something to hide” (e.g., journalists, activists, dissidents), or “were up to no good” (e.g., cybercriminals, terrorists). These inappropriate mental models and misperceptions partially explain why most users overestimate the protection private mode offers.

To summarize the findings above, many participants found utility in private mode (e.g., online shopping, debugging software). However, our participants’ conceptual understanding of the term “private browsing” negatively influenced their usage of private mode in real life. Many incorrectly believed that private mode could be used to send encrypted email, achieve online anonymity, or simply access a phishing webpage because it “felt safer” to do so.

Security goals of private mode. We aimed to further investigate how participants perceived the security goals of private mode. Thus, we asked participants about the entities that could learn about their private browsing activities, and how.

All, but three participants (P03; P18; P25) who identified as security/privacy experts, did not understand what private mode could and could not achieve (i.e., did not recognize the primary security goal of private browsing). Many participants (19 out of 25) believed that a family member, a friend, or a work colleague would not be able to learn about the websites they visited in private mode “whatsoever” (P01). Ten mentioned that this would only be possible if the entity was “technically-sophisticated.” Only P03, P18, and P25 (as mentioned above) correctly explained that private mode protected against a local attacker *after* exiting private mode.

Several participants (12 out of 25) believed that a browser vendor (e.g., Google) could not learn their private browsing activities, citing the following statement that appeared on most browser disclosures: “[Browser vendor] won’t save your information . . .” Further, seven believed that private mode would hide their browsing activities from the employer, six from the ISP, and six from intelligence services and governments.

As we can see, most participants did not understand the main security goal of private mode, partially explaining why several participants perceived those who used private mode as paranoid or up to no good.

Expectations. We then asked participants what they expected from private mode. Again, 19 expected that anyone who had access to their machine would find no evidence of the websites visited privately. Additionally, ten expected that a private mode that worked properly would not link their browsing activities in private mode to those in public mode. 13 also expected that a private mode would protect them from all types of website tracking and ad targeting. Interestingly, five expected a website visited in private mode would not be able to determine whether the user was currently browsing privately or not.

Although some browsers, such as Brave, have added privacy features to reduce online tracking, no browser meets all participants’ expectations. However, we argue that participants’ expectations were high because they overestimated the benefits of private mode.

Part 3: Participatory Design Study

We aimed to investigate why existing browser disclosures did not communicate the actual benefits and limitations of private browsing. To further test participants’ knowledge of private mode, we asked participants to take a short quiz (see Section 3: ‘Methodology’). Participants performed poorly with an average score of 3.21/7.00. Most participants (21 out of 25) overestimated the benefits of private mode.

We also asked participants to explain the following items that appeared on most browser disclosures: history file, cookies, and temporary files. We found that although all participants correctly described a browsing history file, most participants (21 out of 25) either had not heard of a cookie or a temporary file, or did not feel confident explaining what these items meant (in the context of private browsing). These findings suggest that most participants did not understand the functionality of private browsing, a finding recently echoed by [48]. However, we argue (see Section 5: ‘Discussion’) that users do not need to understand the functionality of private browsing in order to use private mode correctly.

We then gave our participants a 15-minute tutorial, and asked them to take the same quiz again. Participants’ quiz performance significantly improved (mean= 6.31/7.00), which was an indication that participants could use the knowledge they newly acquired to critique existing disclosures and then design new ones. Hence, we asked participants to critique the disclosures of Brave, Firefox, and Chrome.

Private mode. Most participants (20 out of 25) criticized Firefox for describing its private mode as a “private window.” Further, 17 participants pointed out that although both Brave and Chrome named their private mode “Incognito,” they still used the phrase “browse privately” in their browser disclosures (in the first sentence), which participants described as “misleading” (P02; P06—P09; P12; P19—P21; P24).

Moreover, 19 participants were confused about when information (e.g., cookies, search items) about websites visited in private mode would get deleted: after “closing a private tab?” (P03), “closing all tabs?” (P09), “closing a [private] window?” (P11), “closing a session?” (P04; P11; P13; P21), or “shutting down a browser?” (P09; P14; P17; P20; P21; P22; P24). Also, five participants questioned whether or not one private session could be shared across multiple windows or tabs.

We also asked participants to suggest a new name for private mode, if any. All participants came up with random names: “non-private,” “everything but private,” “insecure,” “random mode,” and “useless.” Although all participants agreed that the term “private browsing” was misleading, there was no clear winner among the names they suggested.

Primary security goal. Most participants (21 out of 25) pointed out that none of the three disclosures explained the primary security goal of private browsing. Seven participants pointed out that although the Chrome disclosure said that “[a user’s] private browsing activity will be hidden from users sharing the same device,” it did not explain that a user of the device could easily monitor other users’ activities by infecting the device with a malware.

Several participants (17 out of 25) also mentioned that browser disclosures should have mentioned all types of attackers that could violate the security goal of private browsing. They reported that all critiqued disclosures mentioned a subset of all possible attackers (i.e., not the complete set).

Private browsing functionality. 16 out of 25 criticized the use of the following statement by all three disclosures: “[vendor] will save/won’t save the following information.” Participants explained that the statement implied the vendor would not save information on its servers after exiting private mode. Yet, the true meaning of the statement is that the vendor will *only* delete private browsing-related information from the user’s local device, and not necessarily from the vendor’s servers.

Further, 22 out of 25 suggested that the technical explanation of private browsing functionality (e.g., whether cookies would be stored or deleted) should have been hid or deferred until the primary security goal was explained in detail, which none of the disclosures critiqued did. Participants mentioned that browser disclosures should have explained (in bullet points) **what** protection private mode actually offered (protecting against a local adversary). Yet, disclosures described **how** this protection was achieved (e.g., by deleting cookies), without explaining **what** protection the mode offered.

Tracking protection. 12 out 25 participants mentioned that existing browser disclosures should have made it clear that protecting against website tracking was *not* a security goal of private mode. Five participants argued that Brave had been working on reducing online tracking as a browser feature, and not as a private mode feature.

Further, four participants argued most browser vendors did not have the incentive to implement a private browsing mode that delivered the level of privacy expected by consumers, mainly because most web browsers (e.g., Chrome, Internet Explorer) were owned by companies (e.g., Google, Microsoft) that relied on targeting users with advertisements to generate revenue. Hence, participants explained that disclosures should not have used the term “tracking protection” – without explaining what the term meant – to advertise the use of private mode.

Chrome performed better. Many participants (18 out of 25) perceived the Chrome browser disclosure as *relatively* more informative than the disclosures of Brave and Firefox, as it used a list of bullet points to describe both private browsing functionality and attackers. In contrast, nine participants reported that the Brave and Firefox disclosures gave them the false sense that private mode aimed to protect against website tracking and ad targeting, increasing their expectations of the protection offered by private mode beyond reality. Also, eight participants mentioned they would use the private mode of Brave and Firefox to perform sensitive browsing activities (before they were given our tutorial); they said they were influenced by the strong statement in the Brave disclosure: “Private tabs ... always *vanish* when the browser is closed,” and the use of the shield icon by Firefox. Participants explained that the statement and the shield were misleading, and did not communicate the actual benefits of private mode.

Finally, we asked our participants to propose new disclosure designs to better communicate the actual protection of private mode. We discuss the findings in the next section.

DISCUSSION

Our findings show the high-level description of private mode as a “private browsing tab” or a “private browsing window” is not only vague, but also misleading. Users’ mental models and perceptions of the term “private browsing” influence the understanding and usage of private mode in real life. Incorrect or inappropriate mental models – partially derived from this term – could lead users to overestimate the benefits of private mode. For example, some participants used private mode to visit webpages not running HTTPs with a valid TLS certificate, incorrectly believing private mode encrypted Internet traffic. Further, several participants associated private mode with end-to-end encrypted messaging tools, Tor, or VPN.

Additionally, only three participants – who identified as security and privacy experts – correctly and accurately explained the primary security goal of private mode. The vast majority of participants incorrectly believed private mode protected against *any* local attacker, without considering the scenario of a motivated local attacker who could infect a shared machine with a spyware and monitor users’ private browsing activities.

Therefore, it is critical to effectively communicate the actual protection private mode offers. Although users might learn about private mode from peers or online articles, effective browser disclosures remain the vendor’s most reliable channel to communicate information to users. Hence, drawing from the findings of our study and the browser disclosure designs our participants proposed, we distill the following recommendations to improve the design of disclosures:

Explain the primary security goal. As most participants pointed out, none of the browser disclosures they critiqued explained the primary security goal of private mode. Although the Chrome disclosure (version 76.0.3809.100) had the following statement: “Other people who use this device won’t see your activity,” it did not explain that a malicious user of “this device” could monitor the private browsing activities of other users of the same device through a spyware or a key-logger. Thus, disclosures must explain that private mode only protects against an entity that takes control of the user’s machine *after* the user exits private mode.

Explain when information will get deleted. Several participants pointed out that the browser disclosures of both Chrome and Firefox did not explain when information (e.g., browsing history, cookies) about the websites visited in private mode would get deleted. Further, some participants mentioned that although the Brave disclosure (version 0.56.15) had the following statement: “[Information (e.g., cookies)] always vanish when the browser is closed,” it did not clearly communicate the actual functionality of private browsing: information related to a specific private browsing session gets deleted *after* the user terminates that session. Thus, browser designers should better communicate *when* information related to a private browsing session will get deleted (e.g., “when closing a private tab, window, or session.” (P02), “when closing the browser” (P09)).

Explain where information will get deleted from. All three disclosures our participants critiqued had the following statement: “Brave; Chrome; Firefox will not save the following information: browsing history, cookies, ...” Several participants argued that this statement was misleading because it implied information related to a private browsing session would not be stored by the browser vendor. Browser designers should consider rewriting the statement to capture the correct intended meaning: information about websites visited in private mode will not be *locally* stored on the user’s device.

Explain the threat model. Private browsing does not hide activities performed in private mode from motivated local attackers, web attackers, employers, ISPs, browser vendors, and governments. All three critiqued browser disclosures mentioned only a subset of these attackers. Further, several participants mentioned that disclosures needed to clearly describe the entities it protected, and did not protect, against *before* explaining the detailed functionality of private mode, as we explain next.

Hide or defer the explanation of functionality. All three disclosures mentioned different types of files (e.g., browsing history file, cookies, temporary files) that would get deleted after exiting private mode. However, the vast majority of participants did not feel confident explaining what these files meant. Further, several participants preferred that disclosures hid (x13) the explanation of the functionality of private mode or deferred (x9) it until its threat model was described; none of the disclosures our participants critiqued did so.

Notify users when authenticated. We found all participants used private mode while being authenticated to online services, incorrectly thinking their search history would get deleted after exiting the mode. Several participants noted they would like to have a mechanism that would detect when they had started browsing in private mode while being logged into a service.

Avoid using uncertain or misleading words. The Chrome disclosure had the following statement: “Your activity **might** still be visible to [the websites you visit, your employer, etc.]” According to many participants, the use of the word “might” could lead users to incorrectly believe that private mode could protect against, for example, website tracking.

Further, the Brave disclosure stated the following: “Private tabs ... always *vanish* when the browser is closed.” However, it did not explain from *where* the information would get deleted, or rather would vanish, from. The use of the word “vanish” led several participants to incorrectly think that information would completely get removed not only from local devices, but also from web servers.

Explain the utility of private mode. Many participants did not necessarily use private mode to visit “embarrassing websites.” They used the mode to login into an online service using another personal account (e.g., logging into Facebook using two different accounts), debug and test software, or purchase a surprise gift for a family member or a friend. Some participants suggested that browser disclosures should promote the utility of private mode: what private mode could be useful for.

Use bullet points and bold fonts. In line with prior work (see Section 2: ‘Related Work’), most participants used bullet points in their disclosure designs to explain the utility of private mode. Our participants also used bold fonts to emphasize important points (mainly, the main goal of private mode).

Rethink the name “private browsing”. As our results show, the term “private browsing” is misleading. Many participants were “shocked” and felt “vulnerable” upon learning the actual benefits and limitations of private mode. Several participants suggested different names for private mode. However, there was no clear winner among the names suggested. We hypothesize that explaining what protection private mode offers would be sufficient, without the need to mention the name “private browsing” or engineer a metaphor for it.

Finally, we encourage browser designers to validate the design guidelines we extracted from the findings of our study. For future work, we are going to design and test different disclosure prototypes using the Experience Sampling Method (ESM). One possible prototype would be to explain the primary security goal of private mode first, followed by a list of bullet points debunking the myths (or misconceptions) that users have about private mode. Firefox has recently added a link to a list of misconceptions about private mode to its disclosure⁶. However, users who click on the link will be directed to a long page explaining *some* of the misperceptions users have. We hypothesize users do not have the time to read the long page and understand what private mode achieves (and does not achieve). Also, the Firefox disclosure does not explain the main security goal of private mode.

LIMITATIONS

Our study has a number of limitations common to all qualitative research studies. First, the quality of qualitative research mainly depends on the interviewer’s individual skills. Therefore, to minimize bias, one researcher, who was trained to conduct interviews and ask questions in an open and neutral way, conducted all interviews.

Second, some participants’ answers tended to be less detailed. However, the interviewer prompted participants to give full answers. Further, the interviewer gave participants a ten-minute break between the second and third parts of the study, to reduce interviewee fatigue [41].

Third, our qualitative work is limited by the size and diversity of our sample. Following recommendations from prior work to interview between 12 and 25 participants [8], we interviewed 25 participants until new codes stopped emerging.

CONCLUSION

We investigated why most users misunderstand the benefits and limitations of private mode. We did so by designing and conducting a three-part study: (1) a usability inspection of private mode using both cognitive walkthrough and heuristic evaluation; (2) a qualitative, interview-based user study; (3) a participatory design study. We recruited 25 demographically-diverse participants, who used or had used in the past private mode, for the second and third parts of the study.

⁶ <https://tinyurl.com/chi2020-4>

REFERENCES

- [1] Lalit Agarwal, Nisheeth Shrivastava, Sharad Jaiswal, and Saurabh Panjwani. 2013. Do Not Embarrass: Re-Examining User Concerns for Online Tracking and Advertising. In *Proc. Symposium On Usable Privacy and Security*.
- [2] Gaurav Aggarwal, Elie Bursztein, Collin Jackson, and Dan Boneh. 2010. An Analysis of Private Browsing Modes in Modern Browsers. In *Proc. USENIX Security Symposium*.
- [3] Devdatta Akhawe and Adrienne Porter Felt. 2013. Alice in Warningland: A Large-Scale Field Study of Browser Security Warning Effectiveness. In *Proc. USENIX Security Symposium*.
- [4] Bonnie Anderson, Tony Vance, Brock Kirwan, David Eargle, and Seth Howard. 2014. Users Aren't Necessarily Lazy: Using NeuroIS to Explain Habituation to Security Warnings. In *Proc. International Conference on Information Systems*.
- [5] Julio Angulo and Martin Ortlieb. 2015. "WTH..!?!?" Experiences, Reactions, and Expectations Related to Online Privacy Panic Situations. In *Proc. Symposium On Usable Privacy and Security*.
- [6] Cristian Bravo-Lillo, Saranga Komanduri, Lorrie Faith Cranor, Robert W Reeder, Manya Sleeper, Julie Downs, and Stuart Schechter. 2013. Your Attention Please: Designing Security-Decision UIs to Make Genuine Risks Harder to Ignore. In *Proc. Symposium On Usable Privacy and Security*.
- [7] Elie Bursztein. 2017. Understanding Why People Use Private Browsing. <https://elie.net/blog/privacy/understanding-how-people-use-private-browsing/> (2017).
- [8] Kathy Charmaz. 2006. *Constructing Grounded Theory: A Practical Guide through Qualitative Analysis*.
- [9] Neil Chou, Robert Ledesma, Yuka Teraguchi, Dan Boneh, and John C Mitchell. 2004. Client-Side Defense Against Web-Based Identity Theft. In *Proc. Network and Distributed System Security Symposium*.
- [10] Jacob Cohen. 1960. A Coefficient of Agreement for Nominal Scales. *Educational and Psychosocial Measurement* (1960).
- [11] Juliet Corbin and Anselm Strauss. 2014. *Basics of Qualitative Research: Techniques and Procedures for Developing Grounded Theory*. Sage Publications Ltd.
- [12] Rachna Dhamija, J Doug Tygar, and Marti Hearst. 2006. Why Phishing Works. In *Proc. Conference on Human Factors in Computing Systems*.
- [13] DuckDuckGo. 2017. A Study on Private Browsing: Consumer Usage, Knowledge, and Thoughts. <https://spreadprivacy.com/is-private-browsing-really-private/>. (2017).
- [14] Serge Egelman, Lorrie Faith Cranor, and Jason Hong. 2008. You've Been Warned: An Empirical Study of the Effectiveness of Web Browser Phishing Warnings. In *Proc. Conference on Human Factors in Computing Systems*.
- [15] Serge Egelman and Stuart Schechter. 2013. The Importance of Being Earnest [In Security Warnings]. In *Proc. International Conference on Financial Cryptography and Data Security*.
- [16] Adrienne Porter Felt, Alex Ainslie, Robert W Reeder, Sunny Consolvo, Somas Thyagaraja, Alan Bettis, Helen Harris, and Jeff Grimes. 2015. Improving SSL Warnings: Comprehension and Adherence. In *Proc. Conference on Human Factors in Computing Systems*.
- [17] Adrienne Porter Felt, Robert W Reeder, Alex Ainslie, Helen Harris, Max Walker, Christopher Thompson, Mustafa Emre Acer, Elisabeth Morant, and Sunny Consolvo. 2016. Rethinking Connection Security Indicators. In *Proc. Symposium On Usable Privacy and Security*.
- [18] Adrienne Porter Felt, Robert W Reeder, Hazim Almuhiemedi, and Sunny Consolvo. 2014. Experimenting At Scale With Google Chrome's SSL Warning. In *Proc. Conference on Human Factors in Computing Systems*.
- [19] Joseph L. Fleiss, Bruce Levin, and Myunghee Cho Paik. 2013. *Statistical Methods for Rates and Proportions*. John Wiley & Sons.
- [20] Susannah Fox. 2005. Adult Content Online. *Pew Internet & American Life Project* (2005).
- [21] Batya Friedman, David Hurley, Daniel C Howe, Edward Felten, and Helen Nissenbaum. 2002. Users' Conceptions of Web Security: A Comparative Study. In *Proc. Conference on Human Factors in Computing Systems*.
- [22] Xianyi Gao, Yulong Yang, Huiqing Fu, Janne Lindqvist, and Yang Wang. 2014. Private Browsing: An Inquiry on Usability and Privacy Protection. In *Proc. Workshop on Privacy in the Electronic Society*.
- [23] Shirley Gaw, Edward W Felten, and Patricia Fernandez-Kelly. 2006. Secrecy, Flagging, and Paranoia: Adoption Criteria in Encrypted E-mail. In *Proc. Conference on Human Factors in Computing Systems*.
- [24] Greg Guest, Arwen Bunce, and Laura Johnson. 2006. How Many Interviews Are Enough? An Experiment with Data Saturation and Variability. *Field Methods* (2006).
- [25] Russell H. Bernard. 2006. *Social Research Methods: Qualitative and Quantitative Approaches*.
- [26] Hana Habib, Jessica Colnago, Vidya Gopalakrishnan, Sarah Pearman, Jeremy Thomas, Alessandro Acquisti, Nicolas Christin, and Lorrie Faith Cranor. 2018. Away From Prying Eyes: Analyzing Usage and Understanding of Private Browsing. In *Proc. Symposium On Usable Privacy and Security*.

- [27] Tasha Hollingsed and David G. Novick. 2007. Usability Inspection Methods after 15 Years of Research and Practice. In *Proc. International Conference on Design of Communication*.
- [28] Ruogu Kang, Laura Dabbish, Nathaniel Fruchter, and Sara Kiesler. 2015. “My Data Just Goes Everywhere:” User Mental Models of the Internet and Implications for Privacy and Security. In *Proc. Symposium On Usable Privacy and Security*.
- [29] Claire-Marie Karat, Robert Campbell, and Tarra Fiegel. 1992. Comparison of Empirical Testing and Walkthrough Methods in User Interface Evaluation. In *Proc. Conference on Human Factors in Computing Systems*.
- [30] Pedro Giovanni Leon, Justin Cranshaw, Lorrie Faith Cranor, Jim Graves, Manoj Hastak, Blase Ur, and Guzi Xu. 2012. What Do Online Behavioral Advertising Disclosures Communicate to Users?. In *Proc. Workshop on Privacy in the Electronic Society*.
- [31] Arunesh Mathur, Jessica Vitak, Arvind Narayanan, and Marshini Chetty. 2018. Characterizing the Use of Browser-Based Blocking Extensions to Prevent Online Tracking. In *Proc. Symposium On Usable Privacy and Security*.
- [32] Mozilla: Blog of Metrics. 2010. Understanding Private Browsing. <https://blog.mozilla.org/metrics/2010/08/23/understanding-private-browsing/>. (2010).
- [33] Jakob Nielsen and Rolf Molich. 1990. Heuristic Evaluation of User Interfaces. In *Proc. Conference on Human Factors in Computing Systems*.
- [34] Yohko Orito, Kiyoshi Murata, and Yasunori Fukuta. 2013. Do Online Privacy Policies and Seals Affect Corporate Trustworthiness and Reputation? *International Review of Information Ethics* (2013).
- [35] Saurabh Panjwani, Nisheeth Shrivastava, Saurabh Shukla, and Sharad Jaiswal. 2013. Understanding the Privacy-Personalization Dilemma for Web Search: A User Perspective. In *Proc. Conference on Human Factors in Computing Systems*.
- [36] Kristen Purcell, Lee Rainie, and Joanna Brenner. 2012. Search Engine Use. *Pew Internet & American Life Project* (2012).
- [37] Emilee J Rader. 2014. Awareness of Behavioral Tracking and Information Privacy Concern in Facebook and Google. In *Proc. Symposium On Usable Privacy and Security*.
- [38] Lee Rainie, Sara Kiesler, Ruogu Kang, Mary Madden, Maeve Duggan, Stephanie Brown, and Laura Dabbish. 2013. Anonymity, Privacy, and Security Online. *Pew Research Center* (2013).
- [39] Stuart Schecter, Rachna Dhamija, Andy Ozment, and Ian Fischer. 2007. The Emperor’s New Security Indicators: An Evaluation of Website Authentication and the Effect of Role Playing on Usability Studies. In *Proc. IEEE Symposium on Security and Privacy*.
- [40] Douglas Schuler and Aki Namioka. 1993. *Participatory Design: Principles and Practices*. CRC Press.
- [41] Clive Seale. 1999. Quality in Qualitative Research. *Qualitative Inquiry* (1999).
- [42] Andreas Sotirakopoulos, Kirstie Hawkey, and Konstantin Beznosov. 2011. On the Challenges in Usable Security Lab Studies: Lessons Learned from Replicating a Study on SSL Warnings. In *Proc. Symposium On Usable Privacy and Security*.
- [43] Joshua Sunshine, Serge Egelman, Hazim Almuhiemedi, Neha Atri, and Lorrie Faith Cranor. 2009. Crying Wolf: An Empirical Study of SSL Warning Effectiveness. In *Proc. USENIX Security Symposium*.
- [44] Joshua Tan, Lujo Bauer, Joseph Bonneau, Lorrie Faith Cranor, Jeremy Thomas, and Blase Ur. 2017. Can Unicorns Help Users Compare Crypto Key Fingerprints?. In *Proc. Conference on Human Factors in Computing Systems*.
- [45] Janice Y Tsai, Serge Egelman, Lorrie Cranor, and Alessandro Acquisti. 2011. The Effect of Online Privacy Information on Purchasing Behavior: An Experimental Study. *Information Systems Research* (2011).
- [46] Joel Weinberger and Adrienne Porter Felt. 2016. A Week to Remember: The Impact of Browser Warning Storage Policies. In *Proc. Symposium On Usable Privacy and Security*.
- [47] Shomir Wilson, Florian Schaub, Rohan Ramanath, Norman Sadeh, Fei Liu, Noah A Smith, and Frederick Liu. 2016. Crowdsourcing Annotations for Websites’ Privacy Policies: Can It Really Work?. In *Proc. World Wide Web Conference*.
- [48] Yuxi Wu, Panya Gupta, Miranda Wei, Yasemin Acar, Sascha Fahl, and Blase Ur. 2018. Your Secrets Are Safe: How Browsers’ Explanations Impact Misconceptions About Private Browsing Mode. In *Proc. World Wide Web Conference*.